

For questions with **circular bubbles**, you may select exactly *one* choice on Examtool.

- Unselected option
- Only one selected option

For questions with **square checkboxes**, you may select *one* or more choices on Examtool.

- You can select
- multiple squares

For questions with a **large box**, you need to write your answer in the text box on Examtool.

There is an appendix at the end of this exam, containing descriptions of all C functions used on this exam.

You have 170 minutes, plus a 10-minute buffer for distractions or technical difficulties, for a total of 180 minutes. There are 12 questions of varying credit (200 points total).

The exam is open note. You can use an unlimited number of handwritten cheat sheets, but you must work alone.

Clarifications will be posted on Examtool.

Q1 *MANDATORY – Honor Code*

(5 points)

Read the following honor code and type your name on Examtool.

I understand that I may not collaborate with anyone else on this exam, or cheat in any way. I am aware of the Berkeley Campus Code of Student Conduct and acknowledge that academic misconduct will be reported to the Center for Student Conduct and may further result in, at minimum, negative points on the exam and a corresponding notch on Nick's Stanley Fubar demolition tool.

Q2 True/false**(30 points)**

Each true/false is worth 2 points.

Q2.1 TRUE or FALSE: TLS with Diffie-Hellman is still vulnerable to man-in-the-middle attackers by intercepting and performing two separate key exchanges with both sides.

 TRUE FALSE

Q2.2 TRUE or FALSE: In TLS, an attacker cannot change the client and server random values R_b and R_s sent at the beginning of the exchange undetected, even though the client and server have not yet agreed on a set of symmetric keys.

 TRUE FALSE

Q2.3 TRUE or FALSE: Referrer validation is a valid defense against reflected XSS attacks.

Clarification during exam: This question has been dropped. All students will receive credit on this question.

 TRUE FALSE

Q2.4 TRUE or FALSE: Hybrid encryption typically uses symmetric encryption to encrypt the plaintext and asymmetric encryption to encrypt a symmetric key.

 TRUE FALSE

Q2.5 TRUE or FALSE: If the first and last node in a Tor circuit learn that they are part of the same circuit and collude with one another, they are able to deanonymize the Tor user, even if the middle node is honest.

 TRUE FALSE

Q2.6 TRUE or FALSE: One advantage of pointer authentication over stack canaries is that pointer authentication is harder to brute-force than stack canaries.

 TRUE FALSE

Q2.7 TRUE or FALSE: While developing a website, the programmer leaves the server password hidden in the website's HTML. This violates Shannon's Maxim.

 TRUE FALSE

Q2.8 TRUE or FALSE: Most CAPTCHAs are set up to distinguish good bots, such as web crawlers, from malicious bots.

 TRUE FALSE

Q2.9 TRUE or FALSE: Cursorjacking is a UI attack that relies on creating a fake cursor that is more prominent or visible than the real cursor.

 TRUE FALSE

Q2.10 TRUE or FALSE: In a program that has stack canaries enabled, the stack canary takes on the same value across multiple executions of the program.

TRUE FALSE

Q2.11 TRUE or FALSE: If Alice encrypts a message with AES-CBC, but instead of using completely random IVs, she uses R , $R + 1$, $R + 2$, and so on, where R is a random value that she chose once, this scheme is IND-CPA secure.

TRUE FALSE

Q2.12 TRUE or FALSE: DNSSEC uses TLS to protect against attacks when transmitting packets.

TRUE FALSE

Q2.13 TRUE or FALSE: SHA-256 is a good hash function to use when hashing passwords.

TRUE FALSE

Q2.14 TRUE or FALSE: A firewall does not defend against a trusted party inside the firewall that becomes malicious and attempts to breach other computers within the network.

TRUE FALSE

Q2.15 TRUE or FALSE: One weakness of one-time authentication codes, such as those sent to a user's cell phone number, is that they are still vulnerable to a transient phishing attack.

TRUE FALSE

Q2.16 (0 points) TRUE or FALSE: EvanBot is a real bot.

TRUE FALSE

Q3 Piazza Policy**(18 points)**

Q3.1 (5 points) Which of the following URLs have the same origin as `http://piazza.com/`? Select all that apply.

- (A) `https://piazza.com/`
- (B) `http://piazza.com:614/`
- (C) `http://web.piazza.com/`
- (D) `http://piazza.com/run`
- (E) `http://aplaza.com/`
- (F) None of the above

Q3.2 (3 points) If the script `<script src="http://cs161.org/tracking.js"></script>` is included in `http://piazza.com/`, which of these pages can the script modify?

- (G) `http://piazza.com/`
- (H) `http://bank.com/`
- (I) `http://cs161.org/`
- (J) All of the above
- (K) None of the above
- (L) —

Q3.3 (5 points) In which of the following contexts would the contents of a cookie with the `HttpOnly` flag be sent? Select all that apply.

- (A) A user clicks on a link that sends a request over HTTP to the domain of the cookie
- (B) A user clicks on a link that sends a request over HTTPS to the domain of the cookie
- (C) JavaScript is used to send a request over HTTP to the domain of the cookie
- (D) JavaScript is used to read the value of the cookie and send its contents to a different domain
- (E) A page includes an `` tag loading an image using HTTP from the domain of the cookie
- (F) None of the above

Q3.4 (5 points) `http://evanbot.piazza.com` is setting a cookie. Which of these cookie attributes can `http://evanbot.piazza.com` set (without being rejected by the browser) so that the cookie gets sent on a request to `http://web.piazza.com/pictures`? Select all that apply.

- (G) `domain=evanbot.piazza.com; path=/pictures`
- (H) `domain=piazza.com; path=/pictures`
- (I) `domain=com; path=/pictures`

(J) domain=evanbot.piazza.com

(K) domain=piazza.evanbot.com; path=/pictures

(L) None of the above

Q4 Coffee-Shop Attacks**(17 points)**

Dr. Yang comes to MoonBucks and tries to connect to the network in the coffee shop. Dr. Yang and `http://www.piazza.com` are communicating through TCP. Mallory is an on-path attacker.

Q4.1 (5 points) Which of the following protocols are used when Dr. Yang first connects to the Wi-Fi network and visits `http://www.piazza.com`? Assume any caches are empty. Select all that apply.

- (A) CSRF (C) DNS (or DNSSEC) (E) DHCP
 (B) IP (D) HTTP (F) None of the above

Q4.2 (3 points) Suppose Mallory spoofs a packet with a valid, upcoming sequence number to inject the malicious message into the connection. Would this affect other messages in the connection?

- (G) Yes, because the malicious message replaces some legitimate message
 (H) Yes, because future messages will arrive out of order
 (I) No, because on-path attackers cannot inject packets into a TCP connection
 (J) No, because TCP connections are encrypted
 (K) —
 (L) —

Q4.3 (3 points) To establish a TCP connection, Dr. Yang first sends a SYN packet with $\text{Seq} = 980$ to the server and receives a SYN-ACK packet with $\text{Seq} = 603$; $\text{Ack} = 981$. What packet should Dr. Yang include in the next packet to complete the TCP handshake?

- (A) SYN-ACK packet with $\text{Seq} = 981$; $\text{Ack} = 604$
 (B) SYN-ACK packet with $\text{Seq} = 604$; $\text{Ack} = 981$
 (C) ACK packet with $\text{Seq} = 981$; $\text{Ack} = 604$
 (D) ACK packet with $\text{Seq} = 604$; $\text{Ack} = 981$
 (E) Nothing to send, because the TCP handshake is already finished.
 (F) —

Q4.4 (3 points) Immediately after the TCP handshake, Mallory injects a valid RST packet to the server. Next, Mallory spoofs a SYN packet from Dr. Yang to the server with headers $\text{Seq} = X$. The server responds with a SYN-ACK packet with $\text{Seq} = Y$; $\text{Ack} = X + 1$. What is the destination of this packet?

- (G) Dr. Yang (H) The server

(I) Mallory

(K) —

(J) None of the above

(L) —

Q4.5 (3 points) Which of the following network attackers would be able to perform the same attacks as Mallory?

Clarification during exam: By “perform the same attacks,” we mean “reliably perform the same attacks.”

(A) A MITM attacker between Dr. Yang and the server

(D) None of the above

(B) An off-path attacker

(E) —

(C) All of the above

(F) —

Q5 Dual Asymmetry**(15 points)**

Alice wants to send two messages M_1 and M_2 to Bob, but they do not share a symmetric key.

Clarification during exam: Assume that p is a large prime and that g is a generator mod p , like in ElGamal. Assume that all computations are done modulo p in Scheme A.

Q5.1 (3 points) Scheme A: Bob publishes his public key $B = g^b$. Alice randomly selects r from 0 to $p-2$. Alice then sends the ciphertext $(R, S_1, S_2) = (g^r, M_1 \times B^r, M_2 \times B^{r+1})$.

Select the correct decryption scheme for M_1 :

- | | |
|---|--|
| <input type="radio"/> (A) $R^{-b} \times S_1$ | <input type="radio"/> (D) $B^b \times S_1$ |
| <input type="radio"/> (B) $R^b \times S_1$ | <input type="radio"/> (E) — |
| <input type="radio"/> (C) $B^{-b} \times S_1$ | <input type="radio"/> (F) — |

Q5.2 (3 points) Select the correct decryption scheme for M_2 :

- | | |
|---|--|
| <input type="radio"/> (G) $B^{-1} \times R^{-b} \times S_2$ | <input type="radio"/> (J) $B^{-1} \times R \times S_2$ |
| <input type="radio"/> (H) $B \times R^{-b} \times S_2$ | <input type="radio"/> (K) — |
| <input type="radio"/> (I) $B^{-1} \times R^b \times S_2$ | <input type="radio"/> (L) — |

Q5.3 (4 points) Is Scheme A IND-CPA secure? If it is secure, briefly explain why (1 sentence). If it is not secure, briefly describe how you can learn something about the messages.

Clarification during exam: For Scheme A, in the IND-CPA game, assume that a single plaintext is composed of two parts, M_1 and M_2 .

- | | | |
|--------------------------------------|-----------------------------|-----------------------------|
| <input type="radio"/> (A) Secure | <input type="radio"/> (C) — | <input type="radio"/> (E) — |
| <input type="radio"/> (B) Not secure | <input type="radio"/> (D) — | <input type="radio"/> (F) — |

Q5.4 (5 points) Scheme B: Alice randomly chooses two 128-bit keys K_1 and K_2 . Alice encrypts K_1 and K_2 with Bob's public key using RSA (with OAEP padding) then encrypts both messages with AES-CTR using K_1 and K_2 . The ciphertext is $\text{RSA}(\text{PK}_{\text{Bob}}, K_1 || K_2), \text{Enc}(K_1, M_1), \text{Enc}(K_2, M_2)$.

Which of the following is required for Scheme B to be IND-CPA secure? Select all that apply.

- (G) K_1 and K_2 must be different
- (H) A different IV is used each time in AES-CTR

- (I) M_1 and M_2 must be different messages
- (J) M_1 and M_2 must be a multiple of the AES block size
- (K) M_1 and M_2 must be less than 128 bits long
- (L) None of the above

Q6 Under Pressure

(16 points)

Alice and Bob are communicating large pieces of secret data with each other using an IND-CPA secure encryption scheme, but they are being slowed down by their connection! They decide to use compression to improve the amount of data they can send.

Consider the following properties of compression algorithms:

- Compression algorithms reduce the length of the input.
- High-entropy (random-looking) data may only have its length reduced slightly, or not at all
- Low-entropy (predictable) data can often have its length reduced considerably

Q6.1 (3 points) Alice and Bob consider first encrypting and then compressing their data. They send $\text{compress}(\text{Enc}(K, M))$, where the input to the compression algorithm is the output of the encryption algorithm. Provide **one** reason why this might be a bad idea.

Q6.2 (5 points) Realizing their mistake, Alice and Bob instead decide to compress their data before encrypting it. They send $\text{Enc}(K, \text{compress}(M))$, where the input to the encryption algorithm is the output of the compression algorithm.

TRUE or FALSE: This combination of compress-then-encrypt is IND-CPA secure. If you answer True, briefly justify your answer (no formal proof needed). If you answer False, describe how an adversary could win the IND-CPA game with probability greater than 0.5.

- (G) True (H) False (I) — (J) — (K) — (L) —

For the rest of this question, consider a compress-then-encrypt algorithm with the following properties:

- Assume that Alice and Bob’s messages consist of byte strings.
- When compressing, any chain of m consecutive, identical runs of n bytes (total length mn) are compressed into a single run of length n . Runs consisting of a single byte ($n = 1$) are included.
- Encryption preserves the exact length of the message (no padding is used).

For example, the following sequence of bytes:

Index	0	1	2	3	4	5	6	7	8	9
Byte	0x00	0x11	0x22	0x33	0x44	0x44	0x55	0x66	0x77	0x77

would be compressed to a message of length 8, since there are 2 1-byte runs in positions 4–5 and 8–9.

Similarly, the following sequence of bytes:

Index	0	1	2	3	4	5	6	7	8	9
Byte	0xaa	0xbb	0xaa	0xbb	0xcc	0xdd	0xcc	0xdd	0xee	0xff

would be compressed to 6 bytes, since there are two 2-byte runs in positions 0–3 and two 2-byte runs in positions 4–7.

Q6.3 (5 points) Alice sends the same message $M = \text{secret}$ repeatedly to Bob, using the compression algorithm from above for a compress-then-encrypt scheme.

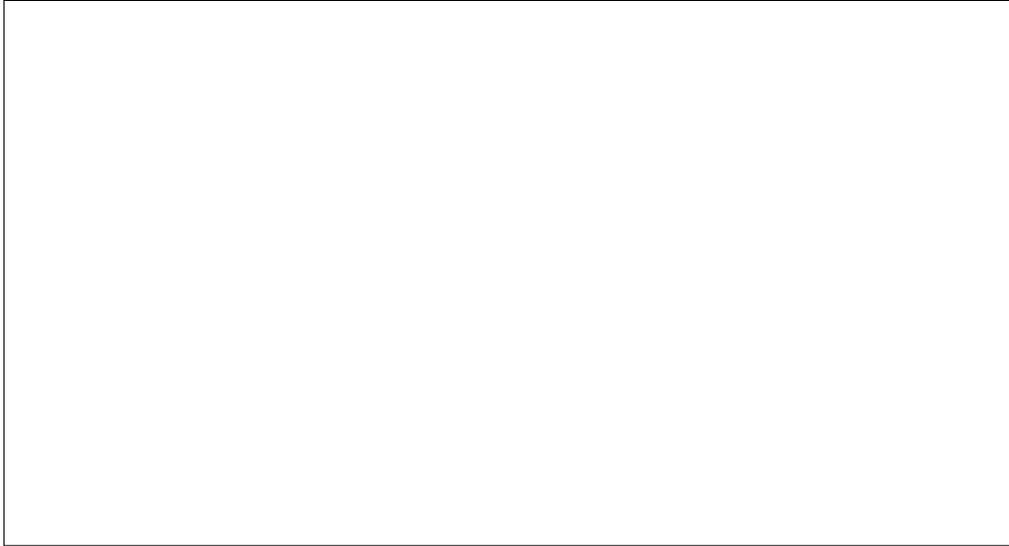
Assume that Mallory doesn’t have complete control over the messages that Alice sends, but Mallory is able to *append* to the sent messages before they are compressed and encrypted.

For example, if Alice repeatedly sends the message $M = \text{secret}$, Mallory can force the compressed and encrypted message to be $M' = \text{secret}||x$, where x is chosen by Mallory. She can do this repeatedly for any value of x that she chooses.

If Mallory can observe all resulting ciphertexts from the compress-then-encrypt algorithm as they are transmitted to Bob, devise a way to recover the **last** byte of secret.

Clarification during exam: Assume that Alice’s messages will never produce ambiguous behavior when being compressed.

Clarification during exam: Assume that the encryption algorithm preserves the exact length of the plaintext.



Q6.4 (3 points) Assume that Mallory has a method for recovering the last byte of the message. If secret is 32 bytes long and there are 256 possible values for a byte, what is the most number of messages Alice has to send (that Mallory can append to) in order to learn the entire secret?

(G) 2^{128}

(I) 32^{256}

(K) 256^{32}

(H) 2^{256}

(J) 256×32

(L) None of the above

Q7 TLS Secret Chaining**(12 points)**

Recall that TLS with RSA does not provide forward secrecy. An attacker who has recorded past connections and then stolen the server's private key can decrypt any past connection.

Consider a modified TLS scheme. RSA is used for the first connection. In future connections, the premaster secret is instead encrypted using the cipher key from the previous connection. Assume this encryption is IND-CPA secure.

The server and client perform n TLS connections. The first connection is numbered C_1 , and the most recent connection is numbered C_n . The attacker records some of these connections and then steals the server's private key. For each set of recorded connections, select all connections the attacker can decrypt.

Clarification during exam: Assume that ranges of connections are inclusive. For example, "All connections between C_1 and C_n " would include C_1 and C_n .

Q7.1 (3 points) The attacker records all connections except C_1 .

- | | |
|--|---|
| <input type="radio"/> (A) Connection C_1 only | <input type="radio"/> (D) All connections between C_1 and C_{n-1} |
| <input type="radio"/> (B) Connections C_1 and C_2 only | <input type="radio"/> (E) All connections between C_1 and C_n |
| <input type="radio"/> (C) All connections except C_1 | <input type="radio"/> (F) None of the above |

Q7.2 (3 points) The attacker records all connections except C_n .

- | | |
|--|---|
| <input type="radio"/> (G) Connection C_1 only | <input type="radio"/> (J) All connections between C_1 and C_{n-1} |
| <input type="radio"/> (H) Connections C_1 and C_2 only | <input type="radio"/> (K) All connections between C_1 and C_n |
| <input type="radio"/> (I) All connections except C_1 | <input type="radio"/> (L) None of the above |

Q7.3 (3 points) The attacker records all connections except C_2 .

- | | |
|--|---|
| <input type="radio"/> (A) Connection C_1 only | <input type="radio"/> (D) All connections between C_1 and C_{n-1} |
| <input type="radio"/> (B) Connections C_1 and C_2 only | <input type="radio"/> (E) All connections between C_1 and C_n |
| <input type="radio"/> (C) All connections except C_1 | <input type="radio"/> (F) None of the above |

Q7.4 (3 points) Suppose we modify the TLS scheme from above. In every connection, the client encrypts their random number R_b with RSA before sending it to the server. (Recall that in regular TLS, R_b is sent with no encryption.)

Does this scheme provide forward secrecy?

- | |
|---|
| <input type="radio"/> (G) Yes, because R_b is needed to generate the symmetric keys |
| <input type="radio"/> (H) Yes, because an attacker can perform a replay attack |

- (I) No, because the attacker can still learn R_b after stealing the private key
- (J) No, because only Diffie-Hellman TLS provides forward secrecy
- (K) —
- (L) —

Q8 *Intrusion Detection Scenarios*

(12 points)

For each scenario below, select the best detector or detection method for the attack.

Q8.1 (3 points) The attacker constructs a path traversal attack with URL escaping: %2e%2e%2f%2e%2e%2f.

- (A) NIDS, because of interpretation issues
- (B) NIDS, because of cost
- (C) HIDS, because of interpretation issues
- (D) HIDS, because of cost
- (E) —
- (F) —

Q8.2 (3 points) The attacker is attacking a large network with hundreds of computers, and a detector must be installed as quickly as possible.

- (G) NIDS, because of interpretation issues
- (H) NIDS, because of cost
- (I) HIDS, because of interpretation issues
- (J) HIDS, because of cost
- (K) —
- (L) —

Q8.3 (3 points) The attacker constructs an attack that is encrypted with HTTPS.

- (A) NIDS, because of interpretation issues
- (B) NIDS, because of cost
- (C) HIDS, because of interpretation issues
- (D) HIDS, because of cost
- (E) —
- (F) —

Q8.4 (3 points) The attacker constructs a buffer overflow attack using shellcode they found online in a database of common attacks.

- (G) Signature-based
- (H) Specification-based
- (I) Anomaly-based
- (J) Behavioral
- (K) —
- (L) —

Q9 To The Moon**(15 points)**

ToTheMoon Bank has just created an online banking system. When a user wants to complete a transfer, they follow these steps:

1. The user logs in by making a POST request with their username and password.
2. The server sets a cookie with `name=auth_user` and `value=$token`, where `$token` is a session token specific to the user's login session.
3. The user initiates a transfer by making a GET request to `https://tothemoonbank.com/transfer?amount=$amount&to=$user`, replacing `$amount` and `$user` with the intended amount and recipient. Transfers use a parameterized SQL query.
4. The server runs the SQL query `SELECT username FROM users WHERE session_token = '$token'`, replacing `token` with the value of the cookie. The server does not use parameterized SQL or any input sanitization.

Q9.1 (4 points) Which of the following attacks are possible in this system? Select all that apply.

- | | |
|--|--|
| <input type="checkbox"/> (A) SQL injection | <input type="checkbox"/> (D) Path traversal attack |
| <input type="checkbox"/> (B) ROP attack | <input type="checkbox"/> (E) None of the above |
| <input type="checkbox"/> (C) CSRF attack | <input type="checkbox"/> (F) — |

Q9.2 (4 points) Mallory is a malicious user with an account on ToTheMoon Bank. Mallory creates a malicious link `https://tothemoonbank.com/transfer?amount=100&to=Mallory`.

Which of the following scenarios would cause Alice to send \$100 to Mallory? Select all that apply.

- (G) Alice clicks on the malicious link when Alice is not logged into the bank
- (H) Alice clicks on the malicious link when Alice is logged into the bank
- (I) Alice visits Mallory's website, which has an `img` tag loading the malicious link, when Alice is not logged into the bank
- (J) Alice visits Mallory's website, which has an `img` tag loading the malicious link, when Alice is logged into the bank
- (K) None of the above
- (L) —

Q9.3 (4 points) Suppose Step 4 is modified. To initiate a transfer, instead of making a GET request, the user makes a POST request to `https://tothemoonbank.com/transfer` with the amount and recipient in the POST body.

Which of the following scenarios would cause Alice to send \$100 to Mallory? Select all that apply.

Clarification during exam: The malicious link in the answer choices should be `https://tothemoonbank.com/transfer?amount=100&to=Mallory`.

- (A) Alice clicks on the malicious link when Alice is not logged into the bank
- (B) Alice clicks on the malicious link when Alice is logged into the bank
- (C) Alice visits Mallory's website, which has an `img` tag loading the malicious link, when Alice is not logged into the bank
- (D) Alice visits Mallory's website, which has an `img` tag loading the malicious link, when Alice is logged into the bank
- (E) None of the above
- (F) —

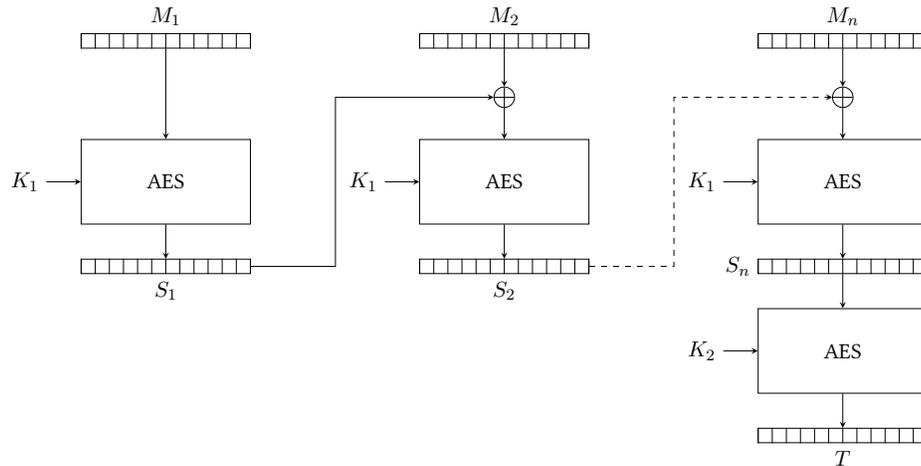
Q9.4 (3 points) Which user inputs might be vulnerable to SQL injection? Select all that apply.

- (G) `amount` parameter
- (H) `to` parameter
- (I) Value of the `auth_user` cookie
- (J) None of the above
- (K) —
- (L) —

Q10 AES-EMAC

(19 points)

Consider AES-EMAC, which is another scheme for generating secure MACs.



Q10.1 (2 points) TRUE or FALSE: Given only T , an attacker can generate a valid MAC for $M||M'$, for an M' of the attacker's choosing.

- TRUE FALSE

Q10.2 (2 points) TRUE or FALSE: Given only T and K_1 , an attacker can generate a valid MAC for $M||M'$, for an M' of the attacker's choosing.

- TRUE FALSE

Q10.3 (2 points) TRUE or FALSE: Given only T , K_1 , and K_2 , an attacker can generate a valid MAC for $M||M'$, for an M' of the attacker's choosing.

- TRUE FALSE

Q10.4 (2 points) TRUE or FALSE: Given the output T and the secret keys K_1 and K_2 , the entire original message M can be reconstructed.

- TRUE FALSE

For the rest of the question, regardless of your answer to the previous parts, assume that the output is $S_1||S_2 \dots ||S_n||T$.

Q10.5 (4 points) Which values are needed to recover the entire original message? Select all that apply.

- (A) S_i for all $1 \leq i \leq n$ (D) K_2
 (B) T (E) None of the above
 (C) K_1 (F) —

Q10.6 (3 points) Which of these equations is correct for calculating a plaintext block M_i given the output and both keys K_1 and K_2 ?

(G) $M_i = \text{Dec}(K_1, S_i) \oplus S_{i-1}$

(J) $M_i = \text{Dec}(K_2, S_i \oplus S_{i-1})$

(H) $M_i = \text{Dec}(K_1, S_i \oplus S_{i-1})$

(K) $M_i = \text{Enc}(K_1, S_i) \oplus S_{i-1}$

(I) $M_i = \text{Dec}(K_2, S_i) \oplus S_{i-1}$

(L) $M_i = \text{Enc}(K_1, S_i \oplus S_{i-1})$

Q10.7 (4 points) Select all true statements about this scheme.

(A) To compute AES-EMAC on a message, the message must first be padded to a multiple of the block size

(B) Encryption can be parallelized

(C) Decryption can be parallelized

(D) AES-EMAC is IND-CPA secure

(E) None of the above

(F) —

Q11 DNS Lookups**(15 points)**

EvanBot performs a lookup for the IP address of `toon.cs161.org`. For each DNS or DNSSEC record, determine which name server sent the record.

Q11.1 (3 points) A type record with the IP address of `toon.cs161.org`

- | | |
|--|---|
| <input type="radio"/> (A) root name server | <input type="radio"/> (D) None of the above |
| <input type="radio"/> (B) <code>.org</code> name server | <input type="radio"/> (E) — |
| <input type="radio"/> (C) <code>cs161.org</code> name server | <input type="radio"/> (F) — |

Q11.2 (3 points) A type record with the IP address of the `cs161.org` name server

- | | |
|--|---|
| <input type="radio"/> (G) root name server | <input type="radio"/> (J) None of the above |
| <input type="radio"/> (H) <code>.org</code> name server | <input type="radio"/> (K) — |
| <input type="radio"/> (I) <code>cs161.org</code> name server | <input type="radio"/> (L) — |

Q11.3 (3 points) A type record with the IP address of `cs161.org` (not the name server)

- | | |
|--|---|
| <input type="radio"/> (A) root name server | <input type="radio"/> (D) None of the above |
| <input type="radio"/> (B) <code>.org</code> name server | <input type="radio"/> (E) — |
| <input type="radio"/> (C) <code>cs161.org</code> name server | <input type="radio"/> (F) — |

Q11.4 (3 points) DNSKEY type record with the public key of the `cs161.org` name server

- | | |
|--|---|
| <input type="radio"/> (G) root name server | <input type="radio"/> (J) None of the above |
| <input type="radio"/> (H) <code>.org</code> name server | <input type="radio"/> (K) — |
| <input type="radio"/> (I) <code>cs161.org</code> name server | <input type="radio"/> (L) — |

Q11.5 (3 points) DS type record with the hash of the `.org` name server's public key

- | | |
|--|---|
| <input type="radio"/> (A) root name server | <input type="radio"/> (D) None of the above |
| <input type="radio"/> (B) <code>.org</code> name server | <input type="radio"/> (E) — |
| <input type="radio"/> (C) <code>cs161.org</code> name server | <input type="radio"/> (F) — |

Q12 Wallet Management**(26 points)**

Consider the following vulnerable C code:

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 struct wallet {
5     char owner[4]; /* 4 bytes. */
6     int amt;      /* 4 bytes. */
7 };
8
9 int main(void) {
10     int wallet_idx = 0;
11     struct wallet wals[8];
12     char buf[16];
13
14     while (1) {
15         /* Get wallet index. */
16         printf("Enter wallet index:\n");
17         fgets(buf, 16, stdin);
18         int wallet_idx = atoi(buf);
19         if (wallet_idx < 0) {
20             /* Exit loop if invalid index. */
21             break;
22         }
23
24         /* Update dollar amount. */
25         printf("Enter dollar amount:\n");
26         fgets(buf, 16, stdin);
27         wals[wallet_idx].amt = atoi(buf);
28
29         /* Read owner. */
30         printf("Enter owner name:\n");
31         gets(wals[wallet_idx].owner);
32     }
33
34     return 0;
35 }
```

Assume you are on a little-endian 32-bit x86 system. Assume that there is no compiler padding or additional saved registers in all subparts.

Clarification during exam: The `atoi` function converts a string to an integer. For example, `atoi("3")` would return the integer value 3.

Q12.1 (4 points) For the first subpart, assume that **no memory safety defenses** are enabled.

Let `SHELLCODE` be a 24-byte malicious shellcode. If the address of `wals` is `0x9ffcad0`, provide an input as a series of Python `print` statements that would cause your program to execute malicious shellcode.

For example, a sequence of non-malicious inputs that updates wallet 0 without exiting the loop would be:

```
print('0')
print('1000')
print('NN')
```

Write your answer in Python 2 syntax (just like in Project 1).

For the remaining parts of this question, assume that **stack canaries are enabled**.

Q12.2 (3 points) Complete the stack diagram with stack canaries enabled. Each row represents 4 bytes.

Parts (2a), (2b), and (2c):

RIP of main
(2a)
(2b)
(2c)
(3a)
(3b)
(3c)
(3d)

- (G) (2a) - SFP of main; (2b) - wallet_idx; (2c) - canary
- (H) (2a) - SFP of main; (2b) - canary; (2c) - wallet_idx
- (I) (2a) - canary; (2b) - SFP of main; (2c) - wallet_idx
- (J) (2a) - canary; (2b) - wallet_idx; (2c) - SFP of main
- (K) (2a) - wallet_idx; (2b) - SFP of main; (2c) - canary
- (L) (2a) - wallet_idx; (2b) - canary; (2c) - SFP of main

Q12.3 (3 points) Parts (3a), (3b), (3c), and (3d):

- (A) (3a) - wals[0].amt; (3b) - wals[0].owner; (3c) - wals[1].amt; (3d) - wals[1].owner
- (B) (3a) - wals[0].owner; (3b) - wals[0].amt; (3c) - wals[1].owner; (3d) - wals[1].amt
- (C) (3a) - wals[7].amt; (3b) - wals[7].owner; (3c) - wals[6].amt; (3d) - wals[6].owner

- (D) (3a) - `wals[7].owner`; (3b) - `wals[7].amt`; (3c) - `wals[6].owner`; (3d) - `wals[6].amt`
- (E) —
- (F) —

Q12.4 (3 points) Describe, briefly, a vulnerability on line 19. Additionally, describe a one-line fix to line 19 that would make this code memory-safe.

Q12.5 (5 points) Let `SHELLCODE` be a 24-byte malicious shellcode. If the address of the RIP of `main` is `0xbfeffc80`, provide an input as a series of Python `print` statements that would cause your program to execute malicious shellcode.

Q12.6 (4 points) Assume that the call to `gets` on line 31 is replaced with `fread(wals[wallet_idx].owner, 1, 4, stdin)`. Recall that `fread` does stop reading at a newline and does not add a NULL terminator.

EvanBot thinks that this code is no longer exploitable because the `gets` function is no longer used. **Assume that your 24-byte shellcode must be written in a contiguous block of memory.** Is EvanBot correct? If you answer Yes, describe why you can no longer exploit this code. If you answer No, describe how you would write your shellcode to a contiguous block of memory.

- (G) Yes
- (H) No
- (I) —
- (J) —
- (K) —
- (L) —

Q12.7 (4 points) This part is independent of the previous subpart, but assume that stack canaries are still enabled. Which of the following actions would individually prevent the attacker from executing malicious shellcode (not necessarily using the exploit from above)?

Clarification during exam: The answer choice “Enabling pointer authentication” has been dropped. All students will receive credit for that answer choice only.

- (A) Enabling non-executable pages in addition to stack canaries
- (B) Enabling pointer authentication

- (C) Replacing the call to `gets` on line 31 with `fgets(wals[wallet_idx].owner, 4, stdin)`
- (D) Swapping the positions of the `owner` and `amt` fields in the `wallet` struct
- (E) None of the above
- (F) —

C Function Definitions

```
int printf(const char *format, ...);
```

printf() produces output according to the format string format.

```
char *gets(char *s);
```

gets() reads a line from stdin into the buffer pointed to by s until either a terminating newline or EOF, which it replaces with a null byte ('\0').

```
char *fgets(char *s, int size, FILE *stream);
```

fgets() reads in at most one less than size characters from stream and stores them into the buffer pointed to by s. Reading stops after an EOF or a newline. If a newline is read, it is stored into the buffer. A terminating null byte ('\0') is stored after the last character in the buffer.

```
size_t fread(void *ptr, size_t size, size_t nmemb, FILE *stream);
```

The function fread() reads nmemb items of data, each size bytes long, from the stream pointed to by stream, storing them at the location given by ptr.

Note that fread() does not add a null byte after input.